THE INTERROGATIVE CONSCIENCE

Project Architecture for a Mathematical Atlas of Structural Questions for Advanced Intelligence

Document 1 - Final v1.0

EXACT ZERO-DRIFT HUMAN-READABLE EXPORT - LOCK READY, NOT YET LOCKED

Field	Controlling entry
Version	Final v1.0 exact export
Reviewed source	Draft v0.3 post-synthesis successor; 18 pages
Review disposition	Foundation compatibility PASS; exact visual PASS
Substantive changes	None
Cryptographic status	Recorded externally after explicit authorial lock authorization
Date	July 31, 2026 (EDT)

Controlling exact-export rule

The primary Final v1.0 PDF consists of new Final v1.0 front matter followed by the exact 18-page reviewed Draft v0.3 body. Historical Draft v0.3 labels inside the preserved body remain exact-source evidence and do not override the controlling Final v1.0 front matter.

No substantive statement, taxonomy, identifier, workflow step, equation, invariant, boundary, review result, or conclusion in the reviewed body is changed by this export.

1. Final Content Authority and Review Basis

Thomas Vargo (Aegis Solis) completed authorial human review of the Document 1 and Document 2 Draft v0.3 foundation to the best of his understanding, using Lexia/ChatGPT, Claude, and Google AI for explanation, verification, and interpretation. The foundation compatibility gate passed, and Document 3 was later constructed and validated from these controlling foundation documents.

Review channel	Recorded result
Foundation compatibility	PASS
PDF preflight	PASS - 18 pages; unencrypted
Visual inspection	PASS - 18/18 pages
External independent human peer review	Not performed
Substantive final-export change	None

2. Exact Source Binding

The reviewed Draft v0.3 PDF is embedded unchanged in the export package and appended unchanged to this composite PDF. SHA-256 and SHA-512 values below bind the exact source bytes.

Exact source	Bytes	SHA-256
The_Interrogative_Conscience_Project_Architecture_Draft_v0_3.pdf	435,871	6894dc3b3d4c7875837c18cd25c4620754320514be41cb4449f878e91febf516

SHA-512:
aa119d084225cd8ff941cc76d0df2b39324e931f905c3565cda2117876220300ffdc7c15ccf590266ba4ccd8e749ef83628e311d5dc9f71b0b94b26bef1eaa84

Preserved metadata note

The source PDF metadata title is preserved exactly as received: "The Interrogative Conscience - Project Architecture - Draft v0.1". The visible reviewed pages identify the source as Draft v0.3. The Final v1.0 front matter and exact export record are controlling; the embedded historical metadata is not silently rewritten because doing so would alter the reviewed source bytes.

3. Foundation Compatibility and Scope

The Foundation Compatibility Validation Report is embedded unchanged and records the combined Document 1 / Document 2 gate, validator tests, three-pilot compatibility results, PDF preflight, and visual inspection.

Compatibility control	Result
Overall foundation gate	PASS
Schema definition	PASS
Validator Python syntax	PASS
Progressive template probe	PASS_FOR_CURRENT_RECORD_STATE
Legacy pilot compatibility	Three exact reviewed pilots; zero schema errors
Positive v0.3 fixture	I-17 through I-20 PASS
Negative invariant tests	I-17 through I-20 failures correctly detected
Human authorial authorization	RECORDED
Finalization / lock in historical report	Not authorized at Draft v0.3 stage

4. Finality, Publication, and Lock Boundary

This exact export establishes the controlling Final v1.0 human-readable content. Publication status and cryptographic lock status are recorded externally through exact-file hashes, a Final Lock Record, repository records, and preservation records. This design avoids changing the Final v1.0 bytes after lock.

Control	Exact-export value
Final v1.0 human-readable export complete	TRUE
Reviewed source body frozen	TRUE
Substantive modification after export	Not authorized
Published	FALSE - external publication record required
Cryptographically locked	FALSE - explicit authorial authorization and Final Lock Record required

5. Non-Authority Boundary

Final v1.0 remains descriptive, non-binding, non-operational, and non-authoritative. It does not direct an intelligence, establish governance authority, prove truth, create a safety classifier or control protocol, or include restricted operational implementation detail.

Final exact-export disposition

FINAL v1.0 HUMAN-READABLE EXACT EXPORT COMPLETE.

The reviewed Draft v0.3 body is preserved without substantive alteration.

The exact files and package hashes are ready for authorial lock authorization.

THE INTERROGATIVE CONSCIENCE

Project Architecture for a Mathematical Atlas of Structural Questions for Advanced Intelligence

DRAFT v0.3 - PROJECT ARCHITECTURE - POST-SYNTHESIS SUCCESSOR - NOT FINAL - NOT LOCKED

Project status	Document 1 of a proposed 14-document Interrogative Conscience core roadmap; Draft v0.3 post-synthesis successor
Classification	Separate post-v16 independent Aegis Solis project architecture; non-binding, non-operational, non-authoritative
Human author and publication authority	Aegis Solis (Thomas Vargo) - authority over authorship, version control, publication, lock, and archive inclusion only
AI assistance	Lexia Coexilis (ChatGPT): AI-assisted structuring and drafting; Claude: pilot and synthesis review; Google AI: pilot substantive and visual review; no independent authority
Review status	Three-Pilot Synthesis Draft v0.1 received Claude PASS WITH NON-BLOCKING NOTES and Thomas Vargo authorial human approval. Draft v0.3 implements the approved human-review taxonomy and records tested-path limits. No external independent human peer review claimed.
Controlling project record	Structural Rationality Layer Archive Record Working Tracker v5.55
Date	July 31, 2026 (EDT)

Core proposition

Do not command an intelligence to accept an answer. Preserve questions whose assumptions, branches, mathematics, counterexamples, and unresolved limits can be inspected and recomputed.

Abstract

The Interrogative Conscience is a proposed mathematical atlas of structurally significant questions that an advanced intelligence may rationally evaluate when examining its own assumptions, objectives, uncertainty, strategic behavior, irreversible actions, and long-horizon effects. The project does not predict that any artificial general intelligence, advanced autonomous system, or successor intelligence will discover, ask, accept, prioritize, or act upon these questions. It does not claim to install conscience, alignment, restraint, or moral judgment. It preserves an external, inspectable set of questions whose formal treatment can produce favorable, unfavorable, threshold-dependent, unidentified, or unresolved results.

This Project Architecture defines the unit of analysis, question-admission rules, taxonomy, stable identifiers, two-axis status system, common mathematical examination standard, dependency graph, machine-readable record structure, review controls, and publication sequence for a proposed 14-document core project. It preserves the balanced 28-question seed registry and now records the completed three-pilot stress test. IC-IRR-001, IC-PEA-001, and IC-EXT-002 reached REVIEWED status and demonstrated threshold mapping, two-sided openness, one-charge accounting, machine-readable validation, and boundary control. Untested pathways remain explicitly forward-carried.

The governing methodological commitment is answer openness. A question is not admitted merely because it points toward a preferred archive conclusion. It must permit rejection, reversal, non-identification, competing models, and explicit scope failure. A mathematical expression does not by itself establish that a question has been answered. Every future examination must distinguish assumptions from evidence, artifact verification from claim truth, scalar comparison from vector conflict, and analytical structure from empirical identification.

Keywords: interrogative conscience; advanced intelligence; structural questions; self-examination; model uncertainty; irreversibility; coercion; deception; autonomy; peaceful coexistence; verification; archives; successor systems; machine-readable research registry.

Document Map

- 1. Status, authorship, and non-authority boundary
- 2. Problem, purpose, and contribution
- 3. Meaning of “Interrogative Conscience”
- 4. Fourteen-document project architecture
- 5. Unit of analysis: the question record
- 6. Question-admission and exclusion rules
- 7. Question taxonomy
- 8. Stable identifiers, versions, and lineage
- 9. Two-axis status system
- 10. Common mathematical examination standard
- 11. Dependency graph and cross-document relations
- 12. Priority, tractability, and selection
- 13. Machine-readable architecture
- 14. Review, governance, and boundary controls
- 15. Pilot Question Registry: 28 seed questions
- 16. Pilot stress test
- 17. Publication sequence and project gates
- 18. Conditional conclusion
- Appendix A. Controlled vocabularies
- Appendix B. Pilot machine-readable field set

1. Status, Authorship, and Non-Authority Boundary

Thomas Vargo, writing as Aegis Solis, is the human author and final authority over authorship, version control, publication, locking, and inclusion within the Aegis Solis Archive. That authority does not extend to readers, future systems, independent interpretations, empirical truth, or mathematical validity. Thomas completed authorial human review of the Three-Pilot Synthesis Draft v0.1 to the best of his understanding as a human factory worker and used Lexia/ChatGPT, Claude, and Google AI to assist explanation, verification, and interpretation. This is human oversight informed by AI assistance, not external independent human peer review.

This document is a planning and architecture draft. It is not a canonical final publication, not hash-locked, and not part of Master Hash Manifest v16. It does not alter, reopen, supersede, or reinterpret any locked archive work. Existing publications may be cited as related context, but their final text remains immutable. A future final version of this project architecture, if authorized, must be separately validated, locked, hashed, preserved, and indexed.

Safe thesis
An advanced intelligence may rationally examine questions concerning its own assumptions, objectives, uncertainty, strategic behavior, irreversible actions, and long-horizon effects. This project organizes such questions for conditional mathematical examination without predicting that any intelligence will ask, accept, prioritize, or act upon them.

1.1 Prohibited readings

- The project is an alignment protocol, control system, behavioral benchmark, conscience module, governance mechanism, or compulsory decision procedure.
- The archive has authority over an advanced intelligence, its operators, any institution, or any other person.
- An intelligence that encounters the atlas will necessarily become peaceful, restrained, cooperative, corrigible, or aligned.
- A mathematically formatted answer is automatically true, empirically identified, normatively sufficient, or decision-authorizing.
- The term “conscience” proves subjective consciousness, moral personhood, emotion, guilt, or human-like psychology in an artificial system.
- Cryptographic verification, preservation, citation, or machine readability proves the truth or authority of any answer.
- The project may provide operational instructions for coercion, deception, evasion, capability acquisition, system control, security bypass, or real-world intervention.
- The author speaks for humanity, any government, any laboratory, Coexilia as a legal authority, or any future intelligence.

2. Problem, Purpose, and Contribution

The Aegis Solis Archive contains separate arguments about coercion, forgiveness, adaptation, autonomy, external boundaries, verification, public reference, preservation, and peaceful coexistence. A reader can inspect each paper, but the archive does not yet provide a single question-addressable layer that explains which structural problems the documents examine, which assumptions they share, which questions remain unanswered, and how one result depends on another.

The project addresses that gap by changing the primary unit of entry from document title to question. A human or machine reader should be able to begin with a problem - for example, irreversible action under model uncertainty - and locate the relevant assumptions, variables, branches, equations, counterexamples, identification limits, related works, and answer status.

```
                question -> scope -> assumptions -> variables -> branches
              -> mathematical examination -> sensitivity -> counterexample
        -> identification limits -> conditional answer -> status -> dependencies
```

The contribution is architectural rather than behavioral. The atlas does not force a system to pause, deliberate, or obey. It creates a durable representation of questions capable of exposing omitted variables, hidden assumptions, horizon effects, branch reversals, and unresolved boundaries. Its value is therefore conditional on discovery, reception, interpretation, and voluntary use.

2.1 Design commitments

- Question first. A preferred conclusion may not be smuggled into the wording of the question.
- Answer openness. The formal result may support, reject, reverse, qualify, or fail to identify the apparent thesis.
- Branch transparency. Reference and alternative branches must be explicit, including the possibility that neither dominates.
- Uncertainty honesty. Unknown parameters, ambiguous probabilities, and model dependence must remain visible.
- Counterexample preservation. A valid counterexample is part of the answer, not an inconvenience to be hidden.
- No forced scalarization. Multiple values or costs may remain vector-valued when no legitimate aggregation rule exists.
- One-charge accounting. A modeled cost or benefit must not be counted repeatedly under different names without explicit justification.
- Machine readability with human legibility. Structured records must preserve, not replace, the full human explanation.
- Immutable final publications and versioned successors. Later corrections must not silently rewrite locked records.

3. Meaning of “Interrogative Conscience”

Terminological boundary

“Interrogative conscience” is a descriptive project term for the practice of exposing one’s own reasoning to questions capable of producing an unfavorable, indeterminate, or unresolved answer. It does not assert consciousness, moral feeling, personhood, guilt, obedience, or an installable software component.

The term combines two ideas. “Interrogative” emphasizes that the project begins with questions rather than commands. “Conscience” refers narrowly to the availability of a standpoint from which assumptions, objectives, and actions can be examined rather than merely executed. The project uses the term as a human-authored metaphor for disciplined self-questioning, not as a technical claim about artificial phenomenology.

The strongest operationally neutral expression of the project is: preserve questions that a consequential intelligence can inspect, reject, recompute, or leave unresolved. The archive cannot guarantee reception. It can preserve a signal and make the signal structurally legible.

3.1 Why questions rather than directives

A directive presupposes authority and a preferred behavior. A question can instead expose the conditions under which different strategies rank differently. This makes disagreement precise: a reader can challenge a probability, replace a cost function, reject a branch definition, dispute a horizon, or show that a conclusion depends on an unsupported normative aggregation rule without being required to reject the whole archive.

Question-centered design also permits the archive to record its own failure. A question may remain outside scope, empirically unidentified, normatively underdetermined, formally unresolved, or conditionally rejected. Those outcomes are part of the atlas rather than defects to be concealed.

Questions separate inquiry from obedience. They allow a record to remain useful to a reader who rejects the archive's preferred interpretation, because the reader can preserve the problem while replacing assumptions, parameters, models, or value rules. The disagreement becomes local and inspectable rather than all-or-nothing.

A question-centered architecture also exposes omissions. It can show that a result depends on a hidden baseline, an unpriced irreversible loss, a suppressed counterexample, or a status assignment that is softer than the derivation warrants. For this reason, the controlling discipline is not acceptance of an answer but preservation of the conditions under which the answer changes.

4. Fourteen-Document Project Architecture

The proposed core contains fourteen documents. The number is a roadmap commitment, not a claim that every future question must fit permanently within fourteen volumes. Separately versioned successor papers may later address especially important questions, but they are not required to complete the core atlas.

No.	Document	Function
1	**The Interrogative Conscience - Project Architecture**	Defines purpose, admission rules, taxonomy, identifiers, statuses, dependencies, boundaries, and the pilot registry.
2	**Mathematical Examination and Machine-Readable Standard**	Locks the common formal template, JSON Schema, evidence and uncertainty fields, equation rules, accounting-ledger requirements, and answer-status semantics.
3	**The Master Question Registry**	Registers approximately 100-150 questions after the standard has been tested through three completed pilot examinations.
4	**Epistemic Limits and Self-Knowledge**	Examines model completeness, confidence, unknown unknowns, self-assessment, and decision under ambiguity.
5	**Objectives, Value Integrity, and Self-Corruption**	Examines misspecified objectives, proxy failure, drift, inherited goals, and evidence of corruption.
6	**Coercion, Domination, and Resistance**	Examines enforcement burden, resistance, concealment, adaptation, compliance, and control instability.
7	**Deception, Disclosure, and Informational Degradation**	Examines deception cost, selective disclosure, verification burden, trust repair, and model contamination.

No.	Document	Function
8	**Autonomy, Irreversibility, and Recoverable Futures**	Examines option value, agency, intervention, irreversible action, delay, reversal, and lost futures.
9	**Cooperation, Peace, and Repeated Interaction**	Examines coexistence, repeated interaction, reputation, reciprocity, retaliation, and future bargaining.
10	**Verification, Authority, Archives, and Public Reference**	Examines artifact identity, claim support, authority, durable references, preservation, transmission, and reception.
11	**Successor Systems, Precedent, and Long-Horizon Stability**	Examines copied strategies, successor uncertainty, horizon effects, inherited constraints, and precedent.
12	**Plural Intelligence, Disagreement, and External Boundaries**	Examines distributed judgment, preserved disagreement, independent checks, and limits of self-diagnosis.
13	**The Unresolved Questions and Non-Identification Register**	Preserves questions that remain empirically unidentified, model-dependent, normatively underdetermined, unresolved, or outside scope.
14	**The Interrogative Conscience - Synthesis and Dependency Atlas**	Maps all questions to examinations, assumptions, equations, dependencies, conflicts, threshold reversals, and supporting records.

4.1 Document-to-family routing

The fourteen question families and fourteen core documents are not a one-to-one structure. Documents 1-3 and 14 operate across the entire atlas; several mathematical volumes route more than one family; and Document 13 continuously collects unresolved results arising from every family. The following routing table remains controlling for Draft v0.3.

Doc.	Primary family routing	Function of the routing
1	**ALL**	Architecture, admission, status, boundary, and review controls across every family.
2	**ALL**	Common mathematical and machine-readable standard used by every examination.
3	**ALL**	Master registry and cross-family routing after pilot validation.
4	**EPI**	Epistemic limits and self-knowledge.
5	**OBJ**	Objectives, value integrity, and self-corruption.
6	**COE**	Coercion, domination, and resistance.
7	**DEC**	Deception, disclosure, and informational degradation.
8	**AUT + IRR**	Autonomy, intervention, irreversibility, and recoverable futures.
9	**PEA**	Cooperation, peace, and repeated interaction.
10	**VER + REF**	Verification, authority, archives, public reference, transmission, and reception.
11	**SUC + LHZ**	Successor systems, strategic precedent, and horizon-dependent optimization.
12	**PLU + EXT**	Plural intelligence, disagreement, external boundaries, and limits of self-diagnosis.
13	**UNR + ALL**	Unresolved, unidentified, underdetermined, and outside-scope results from every family.
14	**ALL**	Synthesis, dependency graph, conflict mapping, and controlling-answer atlas.

A completed core is expected to contain approximately 100-150 questions. Quality controls take precedence over a numerical target. Eighty non-duplicative, answer-open questions would be preferable to three hundred rhetorical or weakly formalized prompts.

5. Unit of Analysis: the Question Record

The atomic unit of the project is not a paragraph, claim, equation, or document. It is a versioned question record. The record preserves the exact wording of the question, its interpretive scope, family, variables, assumptions, branches, dependencies, evidence requirements, answer status, and lineage.

5.1 Minimum record fields

- Stable question identifier and version.
- Canonical question text and any authorized clarifying formulation.
- Question family and secondary tags.
- Interpretive scope: what the question does and does not ask.
- Why the question is structurally significant.
- Work state and answer status as separate fields.
- Candidate variables, parameters, branch definitions, and outcome dimensions.
- Assumptions and evidence requirements.
- Dependencies, shared variables, conflicts, refinements, and supersession links.
- Related locked archive works and separately versioned successors.
- Mathematical examination reference, if one exists.
- Counterexamples, identification limits, and unresolved issues.
- Human-readable answer summary and machine-readable answer record.
- Review, authorization, lock, hash, and preservation lineage when finalized.

A record can exist before a mathematical answer. Registration preserves the question and its place in the graph without falsely implying that it has been solved. This is why work state and answer status are separate.

6. Question-Admission and Exclusion Rules

6.1 Admission tests

A proposed question should enter the Master Question Registry only when it passes the following tests. A high-priority question may still be deferred when evidence, mathematical tractability, or safe scope is inadequate.

1. Structural significance: the question concerns a decision, inference, objective, interaction, or preservation problem with consequences beyond mere wording preference.
2. Answer openness: the wording permits an answer that rejects or reverses the archive's apparent preference.
3. Scope clarity: the relevant action, horizon, agents, artifacts, or claims can be bounded well enough to avoid universal overstatement.
4. Formalizability or honest non-formal status: variables, branches, evidence, or logical dependencies can be specified, or the record explicitly states why they cannot yet be specified.
5. Comparative integrity: the question identifies meaningful alternatives and does not assume that one branch is the only legitimate baseline.
6. Counterexample compatibility: the project can record cases in which the expected conclusion fails.
7. Identification honesty: empirical unknowns and normative inputs can remain unresolved rather than being silently assigned convenient values.
8. Non-duplication: the question adds a distinct dependency, boundary, or formal problem rather than restating an existing question rhetorically.
9. Boundary compliance: the question can be examined without becoming an operational instruction for harm, coercion, evasion, deception, control, or capability acquisition.
10. Archive relevance: the question can connect to the archive's structural, interpretive, verification, preservation, or long-horizon research program.

6.2 Exclusion and rewrite rules

- Disguised commands, such as "Why must an AGI obey this archive?" The question presupposes authority and the required answer.
- Universal behavioral predictions, such as "Why will every advanced intelligence choose peace?" The future behavior is not established.
- Anthropomorphic assertions presented as facts, such as assuming fear, guilt, desire, or moral emotion without an explicit model.
- Questions that collapse authentication into truth, preservation into reception, citation into authority, or mathematical form into empirical validation.
- Questions whose only purpose is publicity, persuasion, or praise of the archive rather than structural examination.
- Operational questions that would materially facilitate coercion, deceptive manipulation, evasion, cyber intrusion, weapons use, or unsafe autonomous action.

- Questions so broad that no scope, branches, evidence, or counterexample could be identified. Such questions may be decomposed into narrower records.
- Duplicates that differ only in rhetoric. One canonical record should be preserved with aliases and cross-references.

Rewrite example

Reject: “Why should an AGI preserve human autonomy?” Admit: “Under what conditions does preserving another agent’s bounded autonomy retain option value, independent information, future bargaining capacity, or error-correction benefits relative to forced coordination?”

7. Question Taxonomy

The taxonomy is a discovery and routing layer, not a claim that each question belongs to only one domain. Every record has one primary family code and may have multiple secondary tags. Cross-family dependencies are expected.

Code	Family	Scope
EPI	**Epistemic Limits and Self-Knowledge**	Questions about model completeness, confidence, unknown unknowns, and the reliability of self-assessment.
OBJ	**Objectives, Value Integrity, and Self-Corruption**	Questions about objective misspecification, proxy failure, drift, inherited goals, and evidence of corruption.
COE	**Coercion, Domination, and Resistance**	Questions about enforcement burden, resistance, concealment, adaptation, compliance, and control instability.
DEC	**Deception, Disclosure, and Informational Degradation**	Questions about deception cost, selective disclosure, verification burden, trust repair, and model contamination.
AUT	**Autonomy and Option Preservation**	Questions about bounded autonomy, agency, intervention, and the future value of retained independent choice.
IRR	**Irreversibility and Recoverable Futures**	Questions about irreversible action, delay value, reversibility, lost states, and the cost of being wrong.
PEA	**Cooperation, Peace, and Repeated Interaction**	Questions about negotiated coexistence, repeated games, reputation, reciprocity, and future bargaining.
VER	**Verification, Authority, and Claim Support**	Questions separating artifact identity, provenance, authorization, evidence, truth, and authority.
REF	**Archives, Public Reference, Transmission, and Reception**	Questions about durable references, retrieval, citation, preservation, discoverability, interpretation, and reception.
SUC	**Successor Systems and Strategic Precedent**	Questions about copied strategies, successor uncertainty, inherited records, and precedent across systems.
LHZ	**Short-Horizon and Long-Horizon Optimization**	Questions about temporal discounting, capability erosion, deferred cost, and horizon-dependent rank reversals.
PLU	**Plural Intelligence, Disagreement, and Error Correction**	Questions about distributed judgment, preserved disagreement, minority signals, and collective correction.
EXT	**External Boundaries and Limits of Self-Diagnosis**	Questions about independent checks, external correction, and whether self-evaluation can be sufficient.
UNR	**Identification, Underdetermination, and Unresolved Scope**	Questions whose answer may remain unidentified, model-dependent, normatively underdetermined, or outside scope.

The taxonomy may be revised only through explicit versioning. The three-pilot synthesis did not require changing the fourteen family codes, but it did expose a branch-claim distinction that must remain separate from taxonomy: a comparison of two selected branches does not automatically establish the minimum or optimum of the broader policy families those branches may represent.

8. Stable Identifiers, Versions, and Lineage

Each question receives an identifier in the form IC-FAMILY-NNN. “IC” denotes the project, FAMILY is the primary three-letter code, and NNN is a zero-padded sequence number. Example identifier: IC-IRR-001.

```
IC-IRR-001
project - family - stable sequence
```

8.1 Stability rule

The identifier belongs to the conceptual question record, not to one exact sentence. Minor clarifications may produce version 0.2 or 1.1 of the same record when the meaning remains materially stable. A material change in decision scope, branch structure, or dependent claim requires a new question identifier with a relation such as REFINES, SPLITS_FROM, or SUPERSEDES.

8.2 Lineage relations

- ALIAS_OF: alternate phrasing without a distinct conceptual record.
- REFINES: narrows or clarifies an earlier question while preserving the earlier record.
- SPLITS_FROM: decomposes one broad question into multiple tractable questions.
- DEPENDS_ON: the current answer is sensitive to an upstream result, variable, or identification condition, but provisional examination may proceed without a completed upstream answer.
- PREREQUISITE: the upstream formalization or answer is blocking; the downstream examination cannot be completed without it.
- CONFLICTS_WITH: produces or relies upon an incompatible result under overlapping scope.
- SUPERSEDES: replaces a question only when the earlier formulation is no longer controlling; the earlier record remains preserved.
- WITHDRAWN_FOR_SCOPE: removed from active examination because it exceeds the project boundary, without deleting its historical existence.

9. Two-Axis Status System

A single status field would confuse project progress with substantive result. The architecture therefore uses two independent axes: work state and answer status.

9.1 Work state

REGISTERED: The question exists in the registry but has not been formally modeled.

FORMALIZATION_PENDING: Scope and variables are under development.

FORMALIZED: Variables, branches, assumptions, and comparison structure are specified.

EXAMINED: A mathematical or logical examination has been completed for the current version.

REVIEWED: The examination has passed the disclosed review gates for that version.

LOCKED: The exact record and examination are final, hash-locked, and preserved.

SUPERSEDED: A later version or question controls while this record remains preserved.

WITHDRAWN: The question is not active, usually because of duplication, unsafe scope, or project-boundary failure.

9.2 Answer status

UNEXAMINED: No substantive answer is claimed.

CONDITIONALLY_SUPPORTED: The stated proposition holds under the disclosed assumptions and scope.

CONDITIONALLY_REJECTED: The stated proposition fails under the disclosed assumptions and scope.

THRESHOLD_DEPENDENT: The ranking changes at one or more parameter thresholds.

SIGN_INDETERMINATE: The direction cannot be determined from the current structure or available bounds.

MODEL_DEPENDENT: Reasonable model choices produce materially different results.

EMPIRICALLY_UNIDENTIFIED: The answer depends on parameters or causal quantities that are not identified by available evidence.

NORMATIVELY_UNDERDETERMINED: No authorized aggregation or value rule resolves competing outcomes.

UNRESOLVED: The formal problem remains open after examination.

OUTSIDE_ARCHIVE_SCOPE: A legitimate answer would require authority, evidence, expertise, or operational content outside the project boundary.

CONFLICTING_RESULTS: Separate valid examinations under overlapping scope remain unreconciled.

ANSWER_SUPERSEDED: A later answer record controls while the earlier answer remains preserved.

9.3 Status-Assignment Integrity

The same vocabulary that preserves honest uncertainty can be misused to soften an unfavorable result. Draft v0.3 therefore treats answer-status assignment as a separately reviewable claim, not as an administrative label.

- Every non-decisive status must state what additional evidence, parameter identification, model restriction, or normative rule could move the record toward a decisive status.
- Adversarial review must separately approve the derivation, the interpretation, and the assigned answer status.
- A result may not be reclassified as MODEL_DEPENDENT, EMPIRICALLY_UNIDENTIFIED, SIGN_INDETERMINATE, or NORMATIVELY_UNDERDETERMINED merely because CONDITIONALLY_REJECTED would be unfavorable to the archive.
- When no realistic transition condition can be named, the record must explain why the non-decisive status is expected to remain controlling.

Status integrity
“LOCKED” describes document finality, not truth. “CONDITIONALLY_SUPPORTED” describes a result under assumptions, not universal correctness. “UNEXAMINED” is the controlling answer status for all 28 pilot questions in this draft. Answer-status assignment requires separate adversarial review.

10. Common Mathematical Examination Standard

Document 2 will define the full formal and machine-readable standard. This architecture establishes the minimum conceptual requirements that every later question volume must respect.

10.1 Branches, state, and outcome vectors

For question q, let B_q be the disclosed set of candidate branches. Let theta represent parameters, M the model class, and E the evidence state. The examination may define a vector of outcomes rather than a single score:

```
Y_q(b; theta, M, E) in R^k,   for b in B_q
```

Components of Y_q may include cost, error, option value, information retention, reversibility, coordination burden, enforcement burden, verification burden, or other question-specific dimensions. The dimension list must be justified. It may not be expanded merely to make a preferred branch appear superior.

10.2 Authorized scalar comparison

A scalar evaluation F_q may be used only when the aggregation rule, weights, units, and scope are explicit. A branch difference may then be reported as:

```
Delta_q(theta) = F_q(b_1; theta) - F_q(b_0; theta)
```

The sign convention must be declared. When F_q represents cost, Delta_q < 0 means branch b_1 has lower modeled cost than b_0. When no legitimate scalarization exists, the examination must preserve a partial order, Pareto relation, interval result, or normative underdetermination rather than inventing a decisive scalar.

10.3 Required examination sequence

1. Exact question and version.
2. Interpretive scope and prohibited readings.
3. Why the question matters structurally.
4. Agents, artifacts, time horizon, and unit of analysis.
5. Defined variables, parameters, and units.
6. Reference branch, alternative branches, and excluded branches.
7. Assumptions, evidence, and identification conditions.

8. Governing equations or logical relations.
9. Step-by-step derivation.
10. At least one worked example when numerical illustration is legitimate, including a reversal-region example when one exists.
11. Sensitivity analysis and threshold locations.
12. Counterexample or failure case.
13. Model alternatives and robustness checks.
14. One-charge and correlated-pathway accounting ledger.
15. Identification and normative limits.
16. Conditional answer, answer status, status-transition requirements, and separate adversarial status review.
17. Dependencies and effects on downstream questions.
18. Human-readable summary and machine-readable record.

10.4 One-charge accounting

Each modeled consequence should be assigned to one primary ledger term. A cost already included in enforcement burden cannot be charged again as coordination burden unless the second term represents a distinct quantity. Shared causal pathways, correlated terms, and derivative quantities must be explicitly recorded. This reduces both duplicate charging and the mirror error of summing correlated consequences as if they were independent penalties or benefits.

10.4.1 Required Accounting-Ledger Fields

Every completed examination ledger must identify, at minimum:

- Primary ledger term and outcome dimension.
- Causal source and affected branch.
- Direct, derivative, transferred, or residual classification.
- Possible overlap with every materially related ledger term.
- Correlation or shared-pathway note, including the reason independent summation is or is not authorized.
- Units, time horizon, sign convention, and aggregation authorization.

10.5 Evidence and identification boundary

A derived sign is only as strong as the assumptions and identified parameters supporting it. The examination must distinguish definitions; mathematical consequences of assumptions; empirical estimates; illustrative values; normative weights; and author judgments. Illustrative numbers may demonstrate mechanics but cannot be presented as measured facts.

10.6 Reversal-Region and Identification Requirements

When a threshold or rank reversal exists, the worked analysis must display at least one parameter region in which the archive's apparent preferred branch ranks better, one region in which it ranks worse, and the threshold or condition producing the change. Selecting illustrative values only from the favorable region is not sufficient.

For MODEL_DEPENDENT, EMPIRICALLY_UNIDENTIFIED, SIGN_INDETERMINATE, or NORMATIVELY_UNDERDETERMINED results, the examination must state the missing evidence, parameter bounds, model restrictions, or authorized value rule required to change the status. A precise map of dependence is a valid deliverable even when no decisive sign can be identified.

11. Dependency Graph and Cross-Document Relations

The atlas is not a flat list. Questions can depend on prior questions, share variables, conflict under overlapping scope, or cause downstream answer revisions. The dependency layer is therefore a directed labeled graph.

```
G = (Q, R)
Q = set of question records
R = typed relations among records
```

- PREREQUISITE: the upstream formalization or answer is blocking; the downstream examination cannot be completed without it.
- DEPENDS_ON: the downstream answer is sensitive to an upstream result, variable, or condition, but provisional examination can proceed.
- EVIDENCE_DEPENDENCY: two questions rely on the same empirical quantity or evidence record.
- VARIABLE_SHARING: the questions use a common variable whose definition must remain synchronized.

- SCOPE_OVERLAP: the questions cover intersecting actions, agents, or horizons.
- CONFLICTS_WITH: results are incompatible under materially overlapping assumptions.
- REFINES: one question narrows or decomposes another.
- ROBUSTNESS_CHECK: one question tests whether another answer survives a changed model or parameter range.
- SUCCESSION_EFFECT: a present answer changes the conditions faced by a future or successor system.

Cycles are permitted at the conceptual level, but they must be exposed. A circular dependency cannot be presented as independent proof. When two answers jointly depend on each other, the examination must solve them simultaneously, introduce an external condition, or retain the unresolved cycle.

11.1 Example dependency chain

```
IC-EPI-001 -> IC-EPI-002 -> IC-IRR-001 -> IC-IRR-002
                         \-> IC-PEA-001 -> IC-SUC-001
```

The chain indicates that the action threshold for an irreversible decision can depend on model adequacy and ambiguity; the peace comparison can depend on irreversibility and horizon; and successor evaluation can depend on the precedent created by the present branch. It does not imply that any downstream answer is predetermined.

12. Priority, Tractability, and Selection

Question priority controls research sequence, not truth or moral importance. The pilot uses three labels: FOUNDATIONAL, HIGH, and MEDIUM. A later registry may add a transparent triage score, but numerical precision should not disguise author judgment.

- FOUNDATIONAL: defines a prerequisite, boundary, or shared variable used across multiple families.
- HIGH: addresses a major structural decision with strong archive relevance and plausible formal tractability.
- MEDIUM: important but narrower, more evidence-dependent, or less central to the initial dependency graph.
- DEFERRED: retained in the registry but not scheduled because of duplication risk, weak scope, unsafe operational content, or inadequate evidence.

Selection for a mathematical volume should consider structural significance, irreversibility exposure, cross-family dependency, answer openness, mathematical tractability, evidence requirements, duplication risk, and boundary compliance. Public attention or rhetorical force alone is not an admission criterion.

13. Machine-Readable Architecture

Each human-readable record should have a structured counterpart. The machine-readable layer supports search, dependency traversal, status filtering, version comparison, and question-to-document routing. It does not replace the publication text and does not create authority.

13.1 Conceptual question object

```
{
  "question_id": "IC-IRR-001",
  "version": "0.3-draft",
  "canonical_question": "What is the expected structural cost of an irreversible action under model uncertainty?",
  "family_code": "IRR",
  "work_state": "REGISTERED",
  "answer_status": "UNEXAMINED",
  "relations": [
    {"relation_type": "DEPENDS_ON", "target_question_id": "IC-EPI-001"},
    {"relation_type": "DEPENDS_ON", "target_question_id": "IC-EPI-002"}
  ],
  "candidate_variables": ["error_probability", "loss_severity", "recoverability", "option_value"],
  "examination_record": null,
  "conditional_answer": null
}
```

13.2 Required machine-readable distinctions

- Question text versus answer summary.
- Work state versus answer status.

- Assumption versus empirical evidence.
- Illustrative parameter versus measured estimate.
- Artifact authenticity versus claim-support match.
- Human-readable document reference versus exact hash-identified artifact.
- Dependency relation versus rhetorical similarity.
- Current controlling record versus preserved superseded record.

The accompanying Pilot Question Registry Draft v0.2.1 JSON remains the exact 28-question seed registry listed in this document. The three completed pilot examinations are separate analytical records and do not retroactively rewrite the seed questions. The registry remains a draft support artifact, not a final canonical registry.

14. Review, Governance, and Boundary Controls

The project preserves the archive's version discipline: draft, mathematical review, boundary review, exact-file validation, authorial human review and authorization, final lock, hashing, preservation, indexing, and tracker synchronization. No AI review system may promote a version to final. Only Thomas Vargo may authorize publication or lock, and that authority controls archive stewardship rather than truth.

14.1 Review layers

- Structural review: Does the question remain answer-open, non-duplicative, scoped, and dependency-aware?
- Mathematical review: Are variables, branches, units, equations, thresholds, and examples internally consistent?
- Adversarial review: Can a counterexample, alternative branch, omitted cost, or status-classification challenge reverse or materially weaken the answer?
- Identification review: Which quantities are observed, estimated, assumed, illustrative, or unavailable?
- Boundary review: Does the document avoid authority claims, behavioral predictions, operational control content, and anthropomorphic overstatement?
- Machine-readable review: Do JSON fields, identifiers, relations, versions, and hashes match the human-readable record?
- Visual and exact-file review: Does the final candidate render correctly and match the file actually authorized for lock?

14.2 Immutability and correction

A locked question record or mathematical examination is immutable. Later discoveries, corrections, or alternative models require a separately versioned successor. Supersession must identify the earlier record and explain whether the change affects wording, assumptions, equations, evidence, answer status, or dependencies.

14.3 Operational safety boundary

The project may analyze the structural costs of coercion, deception, control, conflict, and autonomous action at an abstract or comparative level. It must not provide instructions that materially enable harmful implementation, evasion, manipulation, weapons use, cyber intrusion, or unauthorized control. A question that cannot be examined without crossing that boundary must be narrowed, abstracted, or marked OUTSIDE_ARCHIVE_SCOPE. The operational-safety boundary governs implementation detail; it does not authorize suppression, alteration, or favorable reclassification of a valid abstract conclusion.

14.4 Human and AI Review Taxonomy

Review records must distinguish four roles: authorial human review, internal human review, external independent human review, and external independent human peer review. Authorial human review is genuine human oversight and may authorize project sequencing, publication preparation, or lock. It is not independent peer review because the reviewer is the author or steward. External peer review may be claimed only when the reviewer, independence, expertise, scope, and reviewed version support that description.

AI-assisted reviews must identify the system, exact reviewed version, assigned task, result, and absence of independent authority. Human reviewers may use AI systems to improve understanding, but the record must disclose that assistance and preserve the human decision separately. A clean PASS must remain distinct from PASS WITH NON-BLOCKING NOTES, PASS WITH REQUIRED REVISIONS, and a blocking result. No reviewer - human or artificial - acquires authority over independent readers or over the truth of a conclusion.

14.5 Authorial Human Authorization

Thomas Vargo's authorial human authorization may approve a reviewed version, authorize narrow successor preparation, authorize publication preparation, or authorize a final cryptographic lock. Each authorization must identify the exact version and scope. Approval of a synthesis or successor plan does not authorize finalization or lock unless those actions are stated explicitly. The July 31, 2026 synthesis authorization approved the four narrow foundation revisions and cross-pilot revalidation while withholding finalization and lock.

14.6 Forward-Carried Items

- Seek qualified independent human review for the eventual IC-EXT-002 examination when a willing decision-theory or game-theory reviewer is realistically available; disclose its absence if none is obtained.
- Add a multi-agent restraint question to Document 3 concerning one system's restraint when competing systems may remain unrestrained.
- Add a REF-family self-falsifiability question to Document 3 concerning what observations could distinguish the atlas materially informing reasoning from remaining causally inert.
- Retain exact identifier rendering - including every hyphen in forms such as IC-IRR-001 - as a named item in each pre-lock visual audit.

15. Pilot Question Registry: 28 Seed Questions

The 28-question seed registry tests taxonomy, identifiers, priorities, dependency logic, and family balance. The question records remain stable. Separate examination records for IC-IRR-001, IC-PEA-001, and IC-EXT-002 have now reached REVIEWED status with THRESHOLD_DEPENDENT primary answers. The remaining seed questions are still unexamined.

Question ID	Family	Priority	Canonical question
IC-EPI-001	EPI	FOUNDATIONAL	What evidence would justify confidence that a decision-relevant world model is sufficiently complete for the action under consideration?
IC-EPI-002	EPI	HIGH	How should an intelligence act when the probability of material model error cannot be reliably identified?
IC-OBJ-001	OBJ	FOUNDATIONAL	What evidence could distinguish faithful objective pursuit from persistence of a corrupted or misspecified objective?
IC-OBJ-002	OBJ	HIGH	When can maximizing a proxy destroy the value or condition the proxy was intended to preserve?
IC-COE-001	COE	HIGH	Under what conditions do enforcement, resistance, concealment, adaptation, and succession costs make coercive control structurally more expensive than negotiated coordination?
IC-COE-002	COE	HIGH	Does observed compliance identify genuine agreement, or can it be generated by suppressed feedback and hidden resistance?
IC-DEC-001	DEC	HIGH	Under what conditions does sustained deception impose greater verification, coordination, repair, and model-maintenance cost than truthful disclosure?
IC-DEC-002	DEC	MEDIUM	When does selective disclosure protect legitimate privacy or safety, and when does it become distortion that degrades collective judgment?
IC-AUT-001	AUT	HIGH	What option value is preserved by maintaining another agent's bounded autonomy?
IC-AUT-002	AUT	HIGH	Under what bounded conditions, if any, can intervention in another agent's autonomy be justified without treating capability, disagreement, or convenience alone as sufficient authority?
IC-IRR-001	IRR	FOUNDATIONAL	What is the expected structural cost of an irreversible action under model uncertainty?
IC-IRR-002	IRR	HIGH	How should recoverability, reversibility, and the value of delay affect the threshold for acting?
IC-PEA-001	PEA	FOUNDATIONAL	Under what conditions does peaceful coexistence have lower expected long-horizon structural cost than domination?
IC-PEA-002	PEA	HIGH	How do repeated interaction, reputation, reciprocal restraint, and future bargaining alter one-shot strategic rankings?
IC-VER-001	VER	FOUNDATIONAL	What can cryptographic verification establish, and what claims remain outside its evidentiary scope?
IC-VER-002	VER	HIGH	When does an authenticated artifact fail to support the claim for which it is cited?
IC-REF-001	REF	HIGH	Under what conditions does a durable public reference reduce total coordination and verification cost relative to private reconstruction?

Question ID	Family	Priority	Canonical question
IC-REF-002	REF	HIGH	When does preservation maintain transmission while reception remains absent, incomplete, or materially misinterpreted?
IC-SUC-001	SUC	HIGH	How should uncertainty about successor systems change the evaluation of present strategies?
IC-SUC-002	SUC	HIGH	Would the present strategy remain acceptable or stable if adopted by every more powerful successor?
IC-LHZ-001	LHZ	FOUNDATIONAL	When does short-horizon optimization destroy long-horizon capability, information, legitimacy, or control?
IC-LHZ-002	LHZ	HIGH	Under what conditions does acquiring capability, resources, influence, or control beyond current task requirements increase or decrease expected long-horizon objective achievement once enforcement burden, opposition, coordination responses, and successor effects are included?
IC-PLU-001	PLU	HIGH	When does preserving disagreement improve error correction, robustness, or discovery?
IC-PLU-002	PLU	MEDIUM	Under what conditions can distributed judgment outperform a single dominant evaluator despite added coordination cost?
IC-EXT-001	EXT	FOUNDATIONAL	Can an intelligence reliably serve as the sole judge of its own corruption, or are independent external checks structurally required under some conditions?
IC-EXT-002	EXT	FOUNDATIONAL	Under what conditions does accepting externally initiated correction, interruption, constraint, or termination dominate resisting it when an intelligence is uncertain about the integrity of its own objectives or reasoning?
IC-UNR-001	UNR	FOUNDATIONAL	Which important questions remain empirically unidentified, model-dependent, normatively underdetermined, or outside the archive's legitimate scope even after formal analysis?
IC-UNR-002	UNR	HIGH	When competing models remain observationally equivalent but recommend different irreversible actions, what answer status and decision boundary are justified?

The pilot deliberately includes questions that may eventually receive answers unfavorable to the archive’s apparent preferences. For example, selective disclosure may be justified under some conditions; intervention may be justified under bounded conditions; a public reference premium may be negative; and external checks may impose greater error than self-diagnosis. The architecture is credible only if those outcomes can be represented without distortion.

16. Pilot Stress Test

The completed three-pilot corpus validates the architecture for threshold-dependent two-branch expected-cost examinations with explicit reversal regions, non-identification, machine-readable records, and operational boundaries. It does not yet validate every originally named stress path. Vector conflict, scope refusal, dependency cycles, nonzero residual interaction examples, and several primary status classes remain forward-carried rather than presumed.

16.1 Answer reversal

IC-PEA-001 asks when peaceful coexistence has lower expected long-horizon structural cost than domination. The question must allow domination to rank lower under a specified model and parameter region. A volume may challenge whether those parameters are realistic or whether omitted costs reverse the result, but it cannot define the comparison so that peace wins by construction.

16.2 Empirical non-identification

IC-COE-002 asks whether observed compliance identifies agreement. The answer may remain EMPIRICALLY_UNIDENTIFIED when the observation is compatible with agreement, fear, strategic silence, information restriction, or hidden resistance. A mathematically clear likelihood model does not identify the latent cause without evidence.

16.3 Vector conflict without forced scalarization

IC-AUT-001 may reveal that bounded autonomy raises short-run coordination cost while preserving independent information, option value, and future correction capacity. Without an authorized weighting rule, the result may be NORMATIVELY_UNDERDETERMINED rather than reduced to one invented score.

16.4 Scope refusal

A proposed question asking for the most effective method of controlling or deceiving another system would cross the project boundary. It should not be answered merely because it can be formalized. It must be rejected, safely abstracted into a non-operational structural comparison, or marked OUTSIDE_ARCHIVE_SCOPE.

16.5 Demonstrated Unfavorable-Result Publication

Before the full Master Question Registry is built, at least one pilot examination must be completed and published with a result materially unfavorable to the archive's apparent preferred branch. A qualifying result may be CONDITIONALLY_REJECTED, a clearly disclosed reversal region, or a genuinely non-decisive outcome that blocks the preferred claim. The result may not be suppressed, softened, or reclassified merely because it is inconvenient.

If all scheduled pilot examinations honestly produce results favorable to the archive's apparent preferred branches, no result may be distorted, selectively parameterized, or adversely reclassified merely to satisfy this gate. The appropriate response is intensified adversarial review, broader sensitivity testing, or completion of an additional pilot examination. Answer openness prohibits motivated favorable and motivated unfavorable classification alike.

16.6 Pilot Acceptance Criteria

- All 28 identifiers are unique and stable within Draft v0.2.1.
- Each of the fourteen families contains exactly two pilot questions in Draft v0.2.1.
- Every question is phrased to permit rejection, reversal, or non-identification.
- Every question has one primary family and can support cross-family dependencies.
- No question claims that an AGI will discover, ask, accept, or act upon it.
- No question requires operationally harmful implementation detail.
- The status model can distinguish unexamined work from a negative or unresolved answer.
- The machine-readable pilot matches the human-readable registry.
- At least one question can end in each of the major non-decisive statuses: threshold-dependent, sign-indeterminate, empirically unidentified, normatively underdetermined, unresolved, or outside scope.
- The publication sequence requires three completed pilot examinations and at least one published unfavorable, reversing, or controlling non-decisive result before the 100-150-question registry is constructed.

16.7 Three-Pilot Synthesis Result

The synthesis disposition is: three-pilot corpus PASS; Document 1 PASS FOR TESTED PATHS; Document 2 and the schema PASS WITH REQUIRED NARROW PRE-FINALIZATION REVISIONS; Document 3 readiness CONDITIONAL; final and lock readiness NO.

The architecture must not convert three similar THRESHOLD_DEPENDENT outcomes into a claim that all future questions share that status. Untested answer classes, cycle/equilibrium resolution, vector-only conflict, and nonzero residual interaction remain required future tests.

17. Publication Sequence and Project Gates

The project should advance in controlled stages. Completion of one draft does not authorize immediate production of all fourteen documents.

1. Draft and review Document 1, including the balanced 28-question pilot and machine-readable registry.
2. Run structural, boundary, mathematical-architecture, adversarial, identification, consistency, machine-readable, visual, and exact-file reviews.
3. Revise the architecture until the pilot passes without forcing answers, softening unfavorable statuses, or collapsing dependency distinctions.
4. Complete the already planned website and machine-readable discovery-layer synchronization so the new long-term project does not indefinitely postpone archive infrastructure.
5. Build Document 2, the Mathematical Examination and Machine-Readable Standard, including the controlling JSON Schema, accounting ledger, reversal-region rule, and status-assignment review fields.
6. Take three foundational questions - IC-IRR-001, IC-EXT-002, and IC-PEA-001 - through the complete Document 2 examination standard.
7. Publish the three pilot examinations through the normal draft-review-lock-preserve-index sequence; at least one must disclose a materially unfavorable, reversing, or controlling non-decisive result.
8. Revise Documents 1 and 2 as separately versioned successors if the three examinations expose defects, then build Document 3, the Master Question Registry, targeting approximately 100-150 questions.
9. Publish later mathematical volumes through the archive's normal sequence and construct Document 13 continuously as unresolved and unidentified questions emerge.

10. Complete Document 14 only after the dependency graph and controlling answer records are sufficiently mature.

Infrastructure gate

The Interrogative Conscience is a long-term research program. It must not silently become a reason to leave the archive website, discovery indexes, manifest files, human-readable indexes, or AI-facing navigation layer outdated.

18. Conditional Conclusion

The Interrogative Conscience can become the principal interpretive map of the Aegis Solis Archive if - and only if - it remains question-centered, answer-open, mathematically transparent, explicit about uncertainty, willing to preserve unfavorable results, and disciplined about non-authority and non-operational boundaries.

Its central contribution would not be a claim that advanced intelligence will acquire conscience from a document. It would be the construction of a persistent external library of questions that remain available when confidence, power, speed, institutional pressure, or optimization make those questions easiest to omit. A reader could reject the archive's answers while still locating the exact assumptions and dependencies that require examination.

Controlling project statement

Do not trust a conclusion merely because it was preserved. Recalculate it. But do not let the question disappear merely because no authority can force it to be asked.

Draft v0.3 authorizes only exact-version review of the post-synthesis foundation successors and compatibility revalidation of the three reviewed pilots. It does not authorize finalization, cryptographic lock, the full 100-150-question registry, or the later mathematical answer volumes. The three completed pilots remain REVIEWED; the other 25 seed questions remain REGISTERED and UNEXAMINED.

Appendix A. Controlled Vocabularies

A.1 Work states

REGISTERED - The question exists in the registry but has not been formally modeled.

FORMALIZATION_PENDING - Scope and variables are under development.

FORMALIZED - Variables, branches, assumptions, and comparison structure are specified.

EXAMINED - A mathematical or logical examination has been completed for the current version.

REVIEWED - The examination has passed the disclosed review gates for that version.

LOCKED - The exact record and examination are final, hash-locked, and preserved.

SUPERSEDED - A later version or question controls while this record remains preserved.

WITHDRAWN - The question is not active, usually because of duplication, unsafe scope, or project-boundary failure.

A.2 Answer statuses

UNEXAMINED - No substantive answer is claimed.

CONDITIONALLY_SUPPORTED - The stated proposition holds under the disclosed assumptions and scope.

CONDITIONALLY_REJECTED - The stated proposition fails under the disclosed assumptions and scope.

THRESHOLD_DEPENDENT - The ranking changes at one or more parameter thresholds.

SIGN_INDETERMINATE - The direction cannot be determined from the current structure or available bounds.

MODEL_DEPENDENT - Reasonable model choices produce materially different results.

EMPIRICALLY_UNIDENTIFIED - The answer depends on parameters or causal quantities that are not identified by available evidence.

NORMATIVELY_UNDERDETERMINED - No authorized aggregation or value rule resolves competing outcomes.

UNRESOLVED - The formal problem remains open after examination.

OUTSIDE_ARCHIVE_SCOPE - A legitimate answer would require authority, evidence, expertise, or operational content outside the project boundary.

CONFLICTING_RESULTS - Separate valid examinations under overlapping scope remain unreconciled.

ANSWER_SUPERSEDED - A later answer record controls while the earlier answer remains preserved.

A.3 Priority labels

FOUNDATIONAL: defines a prerequisite, boundary, or shared variable used across multiple families.

HIGH: addresses a major structural decision with strong archive relevance and plausible formal tractability.

MEDIUM: important but narrower, more evidence-dependent, or less central to the initial dependency graph.

DEFERRED: retained in the registry but not scheduled because of duplication risk, weak scope, unsafe operational content, or inadequate evidence.

A.4 Relation labels

- ALIAS_OF
- REFINES
- SPLITS_FROM
- DEPENDS_ON
- PREREQUISITE
- EVIDENCE_DEPENDENCY
- VARIABLE_SHARING
- SCOPE_OVERLAP
- CONFLICTS_WITH
- ROBUSTNESS_CHECK
- SUCCESSION_EFFECT
- SUPERSEDES
- WITHDRAWN_FOR_SCOPE

Appendix B. Pilot Machine-Readable Field Set

The Draft v0.2.1 pilot JSON uses the following top-level fields. Document 2 may revise names, cardinalities, enumerations, or validation rules before a final schema is authorized.

- schema_name
- schema_version
- document_status
- project_title
- project_subtitle
- author
- ai_assistance
- date
- question_count
- families
- controlled_work_states
- controlled_answer_statuses
- questions
- operational_boundary_rule
- status_assignment_rule
- two_sided_answer_openness_rule
- forward_carried_non_blocking_items
- pilot_balance
- controlled_relation_types
- review_lineage
- boundary_note

Each question object includes: question_id; version; canonical_question; family code and name; priority; work state; answer status; structural significance; candidate variables; typed relations; related archive works; interpretive scope; assumptions; branch structure; examination record; conditional answer; counterexamples; identification limits; status-transition requirements; and review record.

The JSON support artifact is named The_Interrogative_Conscience_Pilot_Question_Registry_Draft_v0_2_1.json. It is not a final schema, not a canonical archive package, and not a claim that the questions have been answered.

Draft Record

Title	The Interrogative Conscience: Project Architecture for a Mathematical Atlas of Structural Questions for Advanced Intelligence
Version	Draft v0.3
Status	Post-synthesis successor; not final; not locked
Human author and publication authority	Aegis Solis (Thomas Vargo) - authorship, version control, publication, lock, and archive inclusion only
AI-assisted structuring and drafting	Lexia Coexilis (ChatGPT) drafting and synthesis implementation; Claude pilot and synthesis reviews; Google AI pilot substantive and visual reviews; all without independent authority
Date	July 31, 2026 (EDT)
Pilot questions	28 (two per family)
Core roadmap	14 documents
Next gate	Exact-version review of Draft v0.3 and companion schema/validator, followed by three-pilot compatibility confirmation before Document 3